Current and Future Developments in Networked Collaboration

Mark Handley, USC/ISI

Outline

- **H.323: The myth of transport independence.**
- **Designing for Transport Dependence**
- **Future Directions for Networked Collaboration**

H.3xx Transport Independence

The ITU H.3xx series of standards provide a set of videotelephony standards:

- H.320 for ISDN
- H.321 for POTS
- H.323 for “packet switched networks without service guarantees”
- H.324 for ATM, etc

It is possible to interwork between for example ISDN and Internet clients.

They do this by defining content negotiation mechanisms, common codecs, and so forth.

Gateways are then used to remove one set of transport framing and add another, and MCUs are used for multipoint control.

Transport Independence - the downside

The problem is that these standards have a telephony heritage. Over ISDN, POTS, and (perhaps) ATM, this is not a problem, but over IP the problem and solution spaces are larger and more complex.

- IP multicast supports efficient data distribution to many participants without needing any form of centralised control.
- The Internet is very heterogeneous - a wide range of bandwidths and congestion conditions exist.

Transport Dependencies

- **Design for scalability**
- **Design for heterogeneity**
- **Design for loss**
- **Design for congestion control**

Design for scalability

There are two ways to achieve scalable sessions:

- **Use a hierarchy of servers**
  - Building a hierarchy is hard.
  - Maintaining it in the face of failures is harder.
- **Use multicast**
  - IP multicast sessions can scale to huge numbers of receivers
  - Because multicast is a routing function, it automatically and optimally adapts to failures.

Multicast solutions are not transport independent.

Design for heterogeneity

The internet can be built rapidly, change rapidly, and keep working because it is designed to tolerate heterogeneity.

- TCP is in use on links differing in bandwidth by 8 orders of magnitude!

Layered codecs allow multicast conferencing sessions to scale and to cope well with heterogeneity.

- These codecs split a video stream into multiple additive sub-streams. The more sub-streams you receive, the better the quality.
- Each sub-stream is sent to a different multicast group. Receivers then decide how many groups to subscribe too depending on available network capacity.

Design for loss

Although bandwidth reservation using RSVP may be deployed slowly, it is unlikely to be the ubiquitous solution for preventing packet loss for internet based conferences.

Other enhanced services are also likely to appear, but best-effort service will remain, and will always be cheaper.

Thus conferencing systems should be designed to cope with packet loss gracefully.

Design for loss (2)

Using RTP for transport over IP is a good start

- RTP does application level framing, which copes well with loss

Codecs should also be designed to be loss resilient.

- Motion compensation in H.261 causes packet loss errors to be cumulative.
- You can do “robust H261” but at some expense in compression.
- Codecs designed with this in mind can find a better compromise.

Application level redundancy can also help significantly:

- UCL and INRIA’s robust audio copes well with loss for little additional expense.
- RealAudio stripes data across multiple packets which makes loss much easier to mask.

Design for Congestion Control

There are two ways to use a network:

- reserve what you think you want
- don't reserve, just adapt to what you get

If there is insufficient capacity, reservation failures will make the service too unpredictable to be useful.

If there is sufficient capacity, you don't need a reservation and will often get additional unused bandwidth "for free" if you adapt to what's available.

Design for Congestion Control

If you want to operate using unreserved capacity, you must adapt to congestion signals from the network.

Congestion Control for multicast traffic is currently a very active research area and interesting solutions are starting to emerge.

H.323 gateways between fixed bandwidth pipes and the internet cannot perform congestion control.

- this is because single stream codecs are used, because gateways cannot afford to do transcoding, and because H.320 does not provide for congestion feedback to an H.320 source because an H.320 network cannot be congested!

Transport independence means these streams need to treat the internet as a reserved network.

Transport dependent solutions

- **IP Multicast**
- **Layered Encodings and Receiver Selection**
- **Unreserved network service and congestion control**
- **Robust codecs that cope with loss**
- **Robust transport encodings that protect against loss**
- **Application Level Framing as a design principle**

Where is the future?

If you believe that the future of conferencing is ISDN/ ATM long-haul with IP intranets for local distribution, then perhaps H.323 et al are appropriate standards.

If you believe the future of conferencing will be IP-based, then we can afford solutions that are more tailored to the Internet environment.

Perhaps the Internet cannot afford untailored solutions to be widely deployed?

Future Directions for Networked Collaboration

We have to move significantly beyond video telephony. For five years I've been using the MBone to hold meetings with researchers we've been collaborating with. Some reflections:

- audio quality is still a hard problem due mostly to analog, mike level and silence suppression problems, and causes most disruption.
- the single most useful tool has been the LBL shared whiteboard.
- video is useful in some circumstances, useless in others. It's good for presentations, and for a sense of community, bad for conveying nuances of emotion.
- in general, multimedia conferencing works well for task orientated sessions and badly for management and negotiation.

A few future directions...

Shared collaboration tools are the area for biggest potential improvement in networked collaboration, but are rarely done well.

- T.126 is a half-baked standard
- In general, these tools should probably be allowed to evolve as communities using them grow and change their requirements.
- Reliable multicast is likely to be the basis of many such applications, and we have emerging solutions currently being developed in the research community. Again, congestion control is the key issue.

A few future directions...

Current video tools convey flat images which are useless for negotiation because they fail to convey eye contact (amongst other things).

- To convey eye contact in a group from a single video stream requires three dimensional video encoding so that a virtual scene can be reproduced differently at each receiver depending on who was looking at whom.
- 3D video encoding can be performed using stereo video from two video cameras, although it is currently too expensive to encode in realtime in software.

A few future directions...

Beyond 3D video lies virtual reality...